

AMENDMENTS TO THE CLAIMS

1. (Original) A method for breaking a task into work granules to assign to processes, the method comprising the steps of:
determining how many processes will be used to execute said task;
determining how many granules to divide said task into based on
how many processes will be used to execute said task, and
a range defined by a first threshold and a second threshold;
wherein the first threshold is a minimum number of work granules to assign to each
of the processes that will be used to execute said task;
wherein the second threshold is a maximum number of work granules to assign to
each of the processes that will be used to execute said task; and
dividing said task into a number of work granules that allows each process that will
be used to execute said task to be assigned a number of work granules that
falls within said range.
2. (Original) The method of Claim 1 wherein:
the step of determining how many granules to divide said task into includes the step
of determining a work quantity to equally assign to work granules of said task;
and
the step of dividing said task into a number of work granules includes dividing said
task into work granules that substantially reflect said work quantity.

3. (Original) The method of Claim 2, wherein the step of determining a work quantity includes determining a work quantity that represents an amount of data to access from a database table.
4. (Original) The method of Claim 1, wherein said task entails scanning data that is stored as contiguous sets of data blocks, wherein the method further includes the step of adjusting work assigned to a work granule so that any contiguous set of data blocks scanned by said work granule during execution of the work granule is scanned completely during said execution of the work granule.
5. (Original) The method of Claim 2, wherein each work granule is associated with one or more ranges of ranges of blocks to scan from a database table, said range corresponding to the work quantity assigned to said work granule.
6. (Currently Amended) A computer-readable storage medium carrying~~storing~~ one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 1.
7. (Currently Amended) A computer-readable storage medium carrying~~storing~~ one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 2.
8. (Currently Amended) A computer-readable storage medium carrying~~storing~~ one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 3.

9. (Currently Amended) A computer-readable storage medium carrying~~storing~~ one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 4.
10. (Currently Amended) A computer-readable storage medium carrying~~storing~~ one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 5.
11. (Currently Amended) A computer-readable storage medium carrying~~storing~~ one or more sequences of one or more instructions for managing the assignment of a plurality of work granules to a plurality of processes on a plurality of nodes, the one or more sequences of one or more instructions including instructions which, when executed by one or more processors, cause the one or more processors to perform the steps of:
making a determination that:
a portion of a first work granule of said plurality of work granules has an affinity to a particular node of said plurality of nodes, and
another portion of said first work granule has an affinity to another node of said plurality of nodes;
based on said determination, generating first data that establishes said first work granule as having an affinity to a particular node of said plurality of nodes;
assigning each work granule from the plurality of work granules to a process from said plurality of processes based on said first data;

wherein a portion of each work granule of said plurality of work granules has an affinity for a given node of said plurality of nodes when the portion of said each work granule can be executed more efficiently by said given node relative to another node of said plurality of nodes; and
wherein the step of generating first data includes

generating second data that establishes as not having an affinity for any particular node of said plurality of nodes a second work granule that includes a portion of work which has an affinity for a node from said plurality of nodes, and another portion of work which has an affinity for another node from said plurality of nodes.

12. (Currently Amended) The computer-readable storage medium of claim 11, wherein the step of generating second data includes:
adding said first work granule to a first list that includes work granules established as having an affinity for said particular node; and
adding said second work granule to a second list that includes work granules established as having no affinity for any particular node from said plurality of nodes.
13. (Currently Amended) A computer-readable storage medium ~~carrying~~storing one or more sequences of one or more instructions for managing the assignment of a plurality of work granules to a plurality of processes on a plurality of nodes, the one or more sequences of one or more instructions including instructions which, when

executed by one or more processors, cause the one or more processors to perform the steps of:

making a determination that:

a portion of a first work granule of said plurality of work granules has an affinity to a particular node of said plurality of nodes, and another portion of said first work granule has an affinity to another node of said plurality of nodes;

based on said determination, generating first data that establishes said first work granule as having an affinity to a particular node of said plurality of nodes;

assigning each work granule from the plurality of work granules to a process from said plurality of processes based on said first data;

wherein a portion of each work granule of said plurality of work granules has an affinity for a given node of said plurality of nodes when the portion of said each work granule can be executed more efficiently by said given node relative to another node of said plurality of nodes; and

wherein the step of generating first data includes generating first data that establishes as having an affinity to a particular node from said plurality of nodes a particular work granule that specifies a plurality of database table partitions to scan.

14. (Currently Amended) A computer-readable storage medium ~~carrying~~storing one or more sequences of one or more instructions for managing the assignment of a plurality of work granules to a plurality of processes on a plurality of nodes, the one

or more sequences of one or more instructions including instructions which, when executed by one or more processors, cause the one or more processors to perform the steps of:

making a determination that:

a portion of a first work granule of said plurality of work granules has an affinity to a particular node of said plurality of nodes, and another portion of said first work granule has an affinity to another node of said plurality of nodes;

based on said determination, generating first data that establishes said first work granule as having an affinity to a particular node of said plurality of nodes;

assigning each work granule from the plurality of work granules to a process from said plurality of processes based on said first data;

wherein a portion of each work granule of said plurality of work granules has an affinity for a given node of said plurality of nodes when the portion of said each work granule can be executed more efficiently by said given node relative to another node of said plurality of nodes;

determining that a quantity of certain data is accessed during execution of said first work granule;

determining that a majority of said quantity of said certain data resides at a location local to said first node; and

wherein the step of establishing said first work granule as having an affinity to said particular node is performed in response to determining that a majority of said quantity of third data resides at a location local to said first node.

15. (Currently Amended) A computer-readable storage medium ~~carrying~~storing one or more sequences of one or more instructions for assigning work granules to a plurality of processes on a plurality of nodes, the one or more sequences of one or more instructions including instructions which, when executed by one or more processors, cause the one or more processors to perform the steps of:
dividing a task into a plurality of work granules that includes
a first set of work granules that each define work by specifying one or more ranges of blocks to scan, and
a second set of work granules that each define work by specifying at least two database table partitions of a database table to scan;
generating first data that specifies, for each work granule of said plurality of work granules, whether said each work granule has an affinity for a particular node of said plurality of nodes;
assigning said plurality of work granules to said plurality of processes based on said first data;
wherein said database table is partitioned into said at least two database table partitions by values in one or more columns of said database table; and
wherein the step of generating first data includes generating for each of said first set and said second set:
at least one list that includes work granules having an affinity for a particular node from said plurality of nodes; and
another list that includes work granules having no particular affinity for a particular node.

16. (Currently Amended) The computer-readable storage medium of Claim 15, wherein the step of assigning said plurality of work granules includes assigning to a first process on a first node from said plurality of nodes:
- a work granule from said at least one list when said at least one list includes work granules that have an affinity for said first node and at least one work granule has not been assigned to a process; and
 - a work granule from said another list when said at least one list includes no work granule that has not been assigned to a process.